


UDP, TCP, Circular Buffers, Multi-threaded Programming and the Transmission of Event Data

Richard Riedel*, Piotr Zolnierczuk, Andre Parizzi, Madhan Sundaram
Oak Ridge National Laboratory



Avoiding unpleasant surprises



"Say . . . What's a mountain goat doing way up here in a cloud bank?"



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Event Mode In a Nutshell

- Every neutron detected is time stamped and saved without histogramming.
- Other variables (chopper TDC, magnetic values, temperature values) are also time stamped producing a movie like data set at run end.
- Vetos, gating etc. are done in software..not hardware



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Event Mode Experiments







Figure 4. One second neutron diffraction patterns summed over 100 cycles of the 316LN ss. The zoom-in contour plot shows clearly lattice dilations for (111), (200) directions.



Why should SNS care about higher time resolution dynamic/transient experiment?

- Most intense neutron flux makes instantaneous measurement possible

High flux matters: 1 pulses pattern from VULCAN @900kW.



Thanks to the new power level of SNS and high neutron flux at VULCAN.

Managed by UT-Battelle
for the U.S. Department of Energy

Ke An, kean@ornl.gov

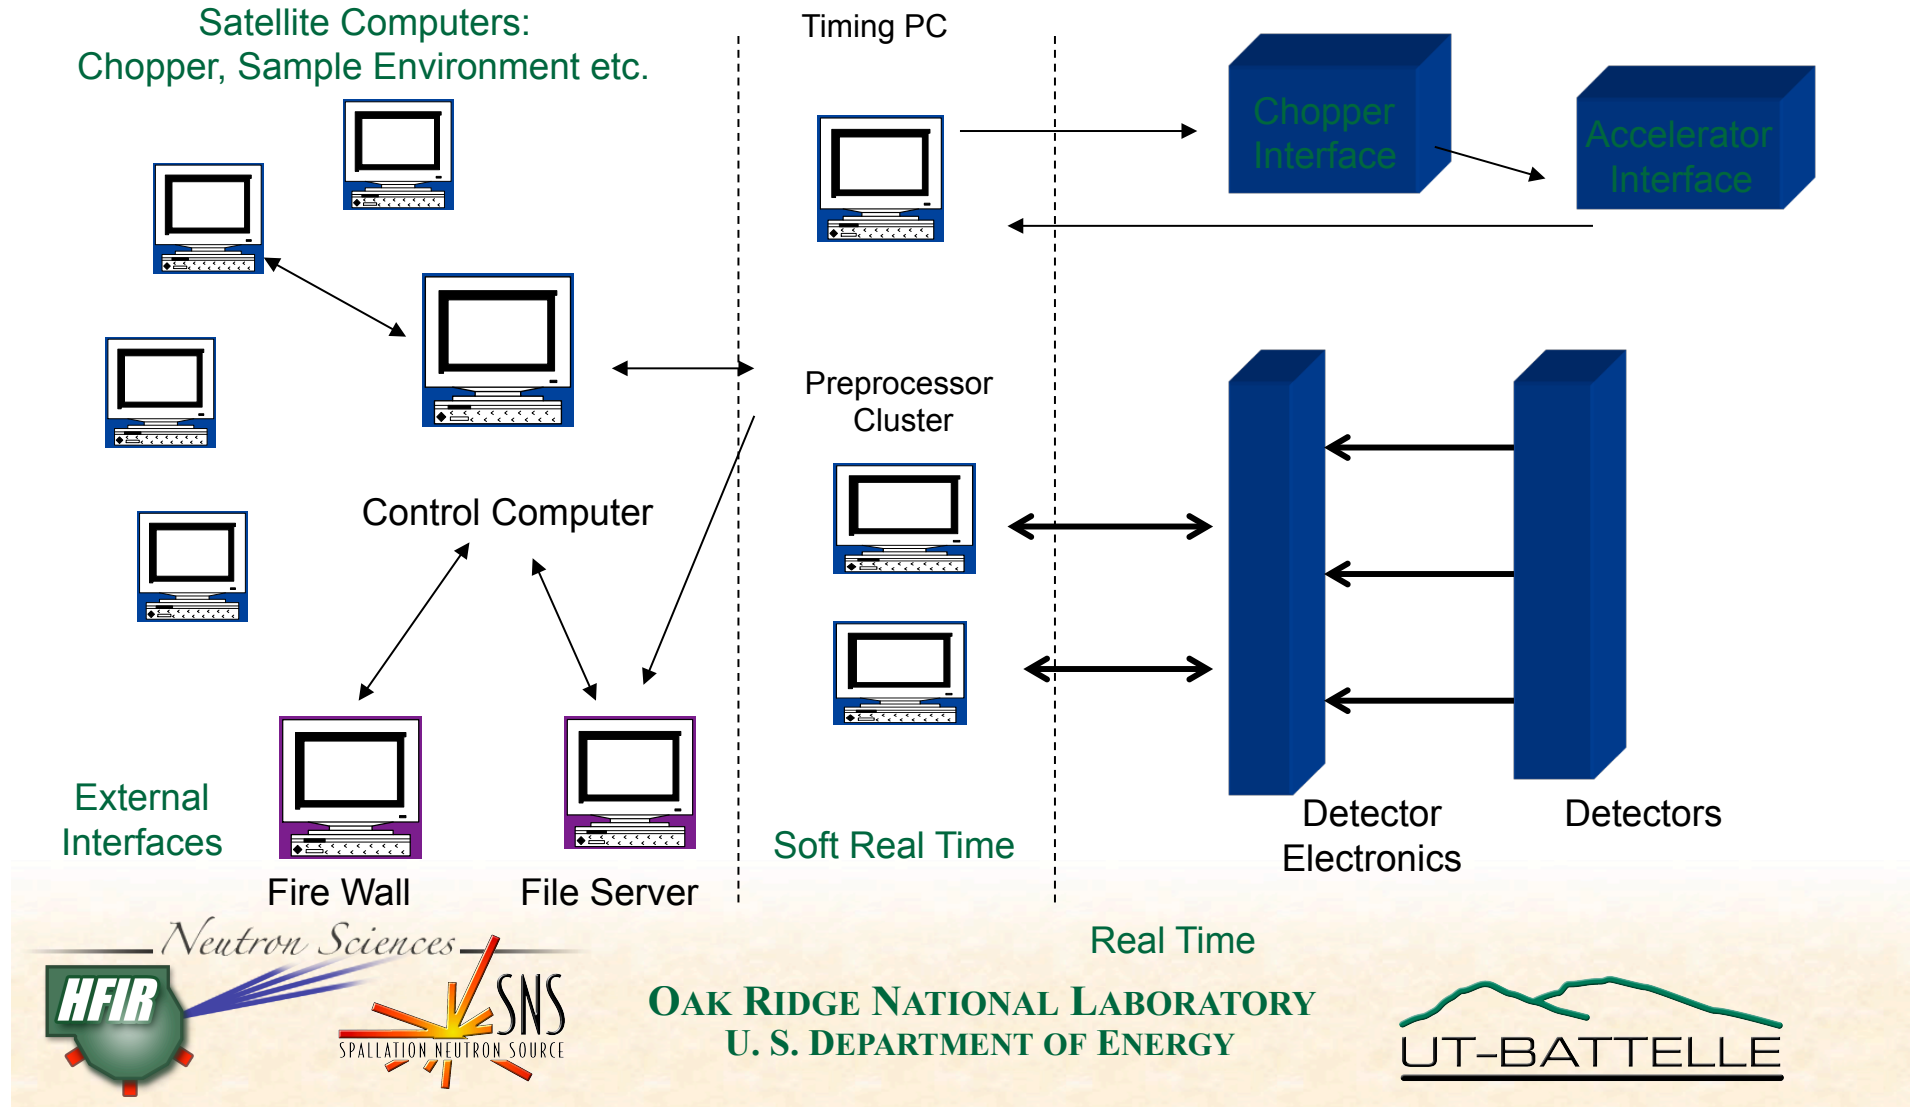


OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

Slides Courtesy of Ke An



DAQ System Diagram



Three+ Links For Communications

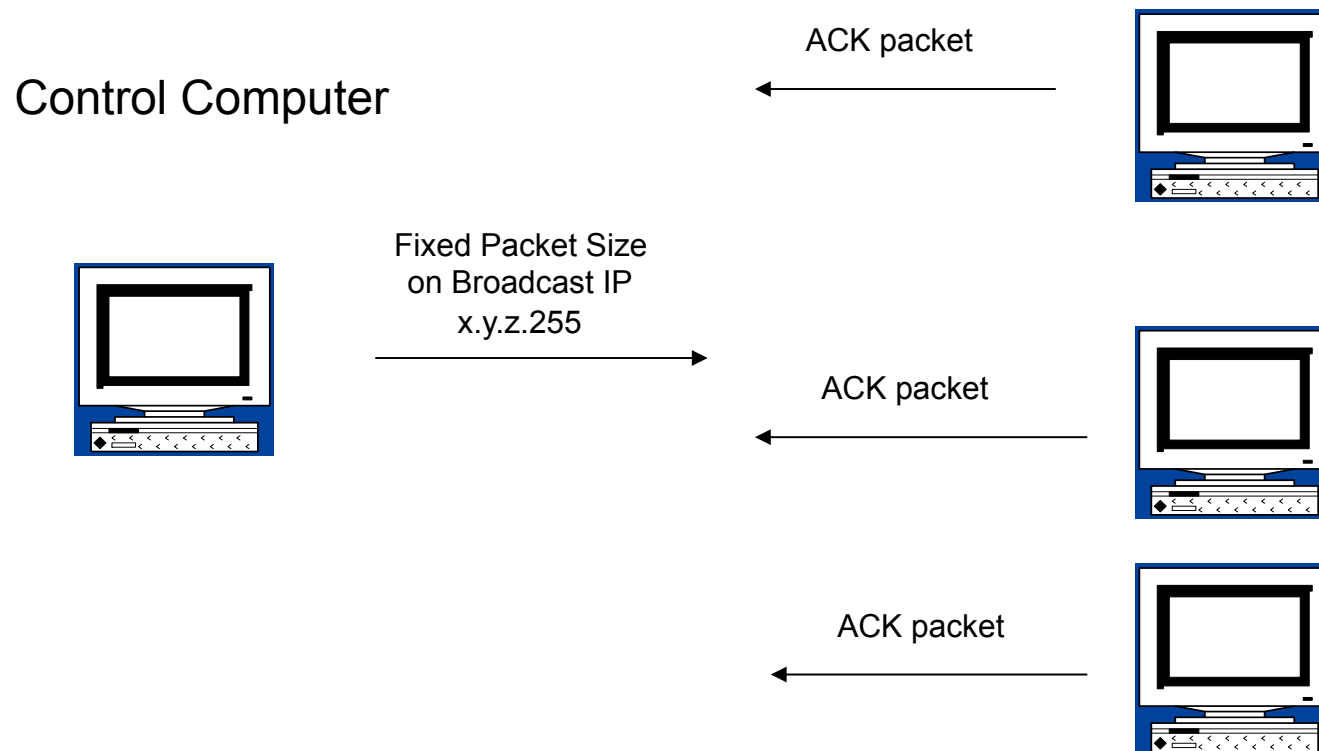
- **Command Link Uses UDP.** Requested action is acknowledge by listeners.
- **Neutron Data Link Uses UDP and TCP.** UDP is broadcast without concern about receipt. TCP is point to point to Data File Server (DFS).
- **Most commands are small (24bytes).**
- **Transfer of other files or data to DFS is done via TCP.** (Never use XP file calls.)
- **Live UDP transfer of meta-data protocols not finalized.**



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Command Link




ACK packets typically separated by < 1msec.



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Transmission Benchmarks



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Issues With Early Implementation.

- **CAsyncSocket** (callback based architecture) missed UDP replies when received < aprox. 1msec apart.
- **Wireshark@** showed UDP packets were always received by card. (True in both directions).
- Class was developed more for convenience than throughput. Also difficult to implement in separate thread.

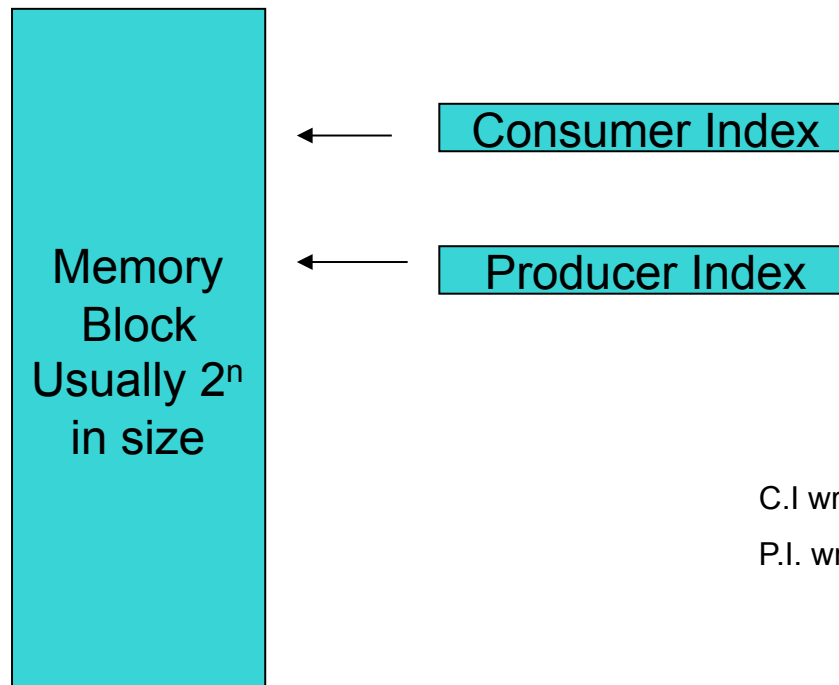


OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Better Approach

- Use base c type Berkley socket functions.
- Use multithreaded listener with circular buffer.



C.I written only by data consuming thread.

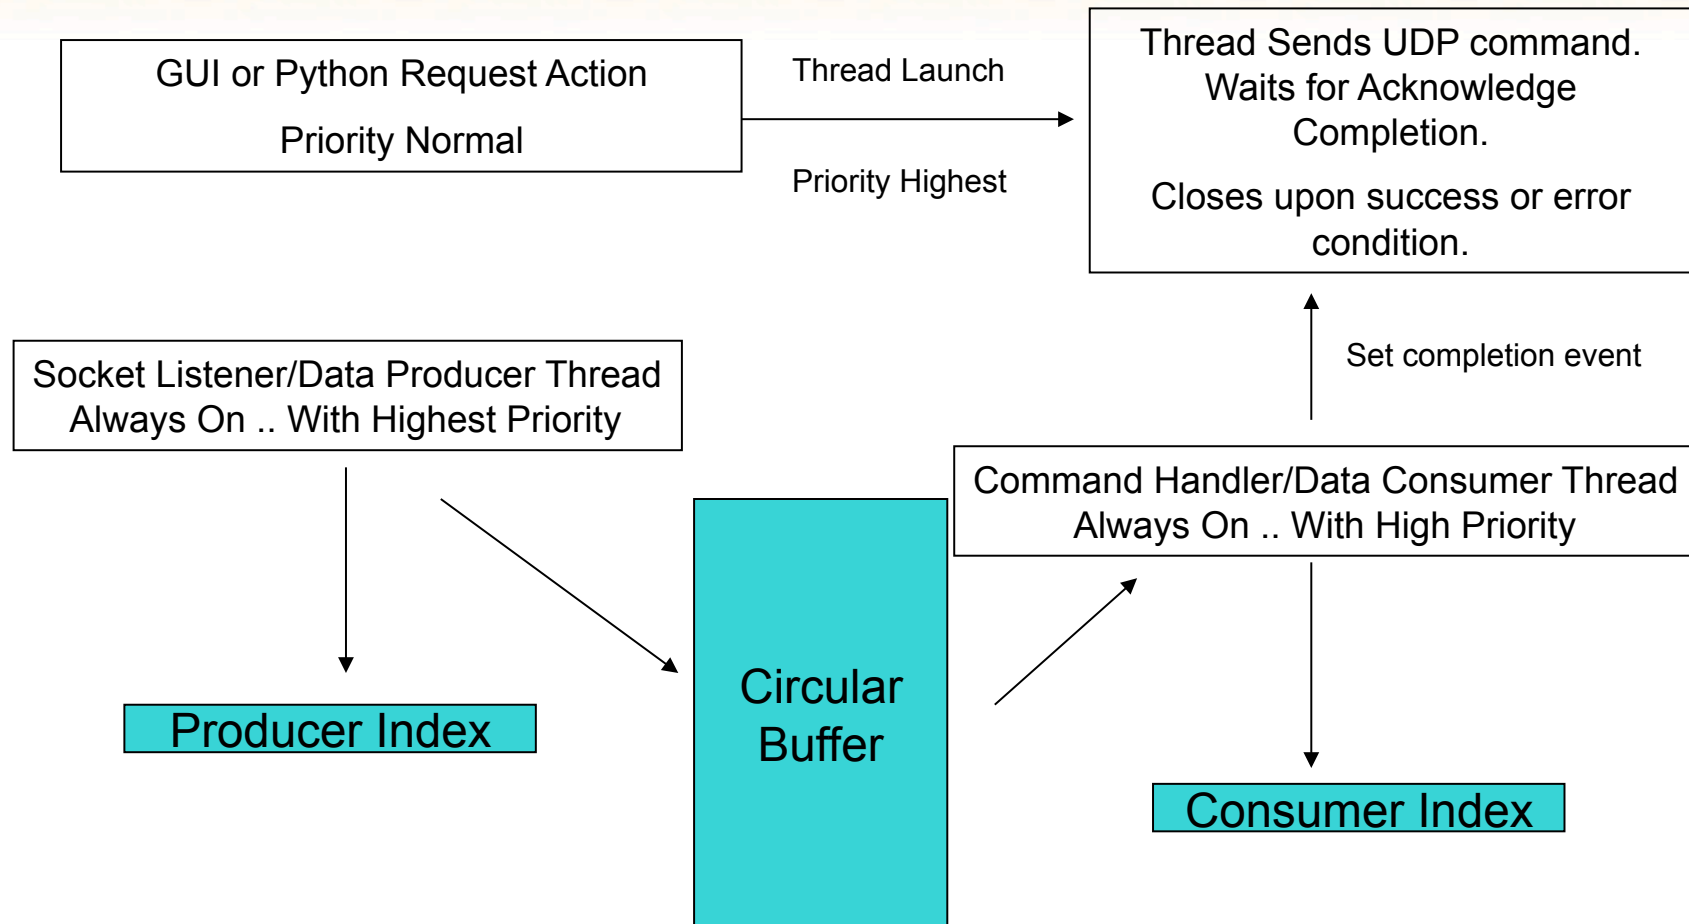
P.I. written only by data producing thread.



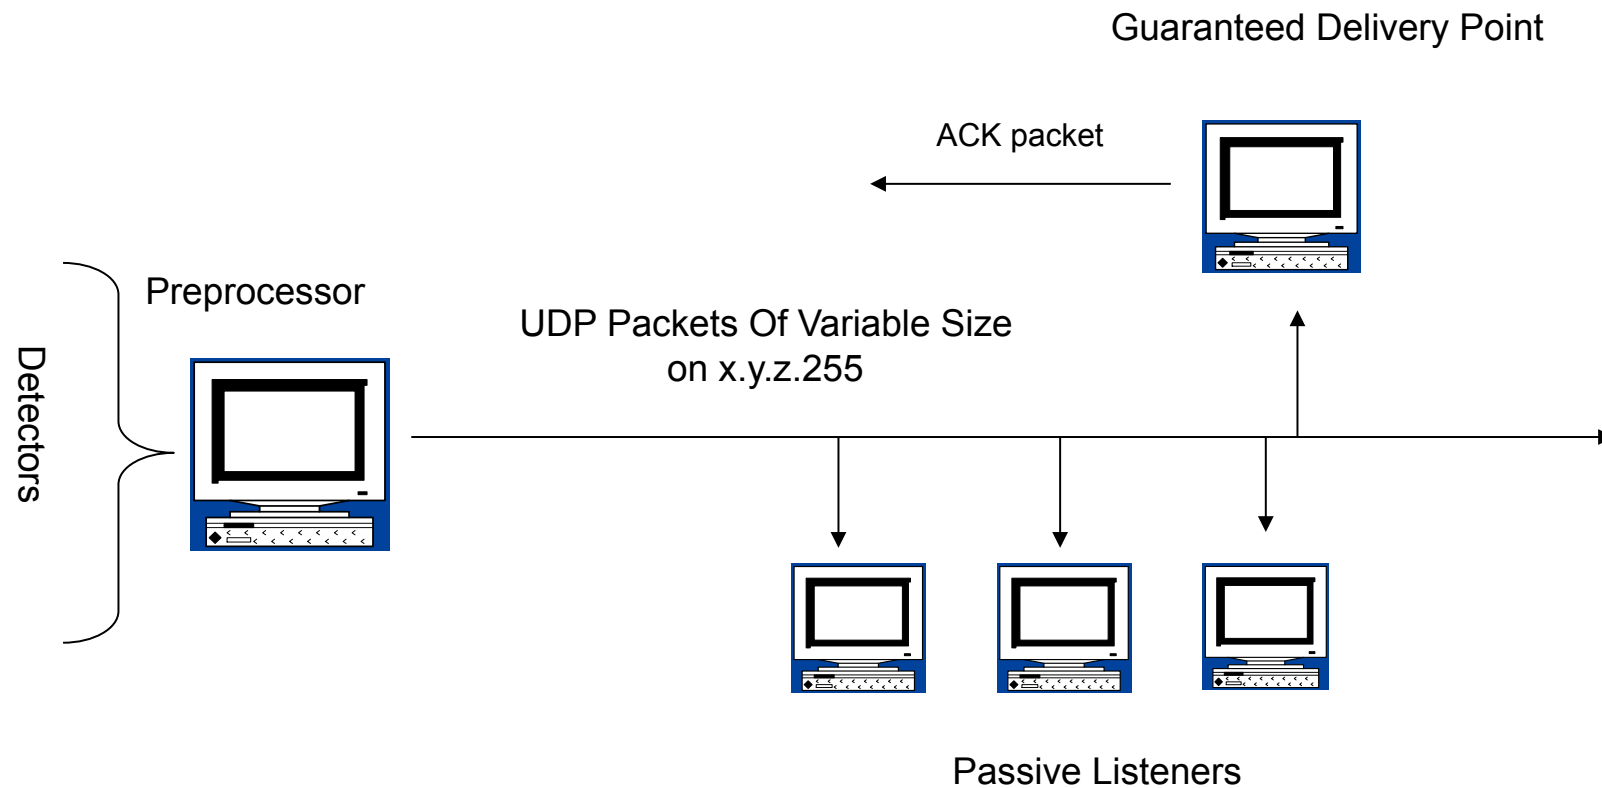
OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Command Link Threading Model (One Producer-One Consumer)

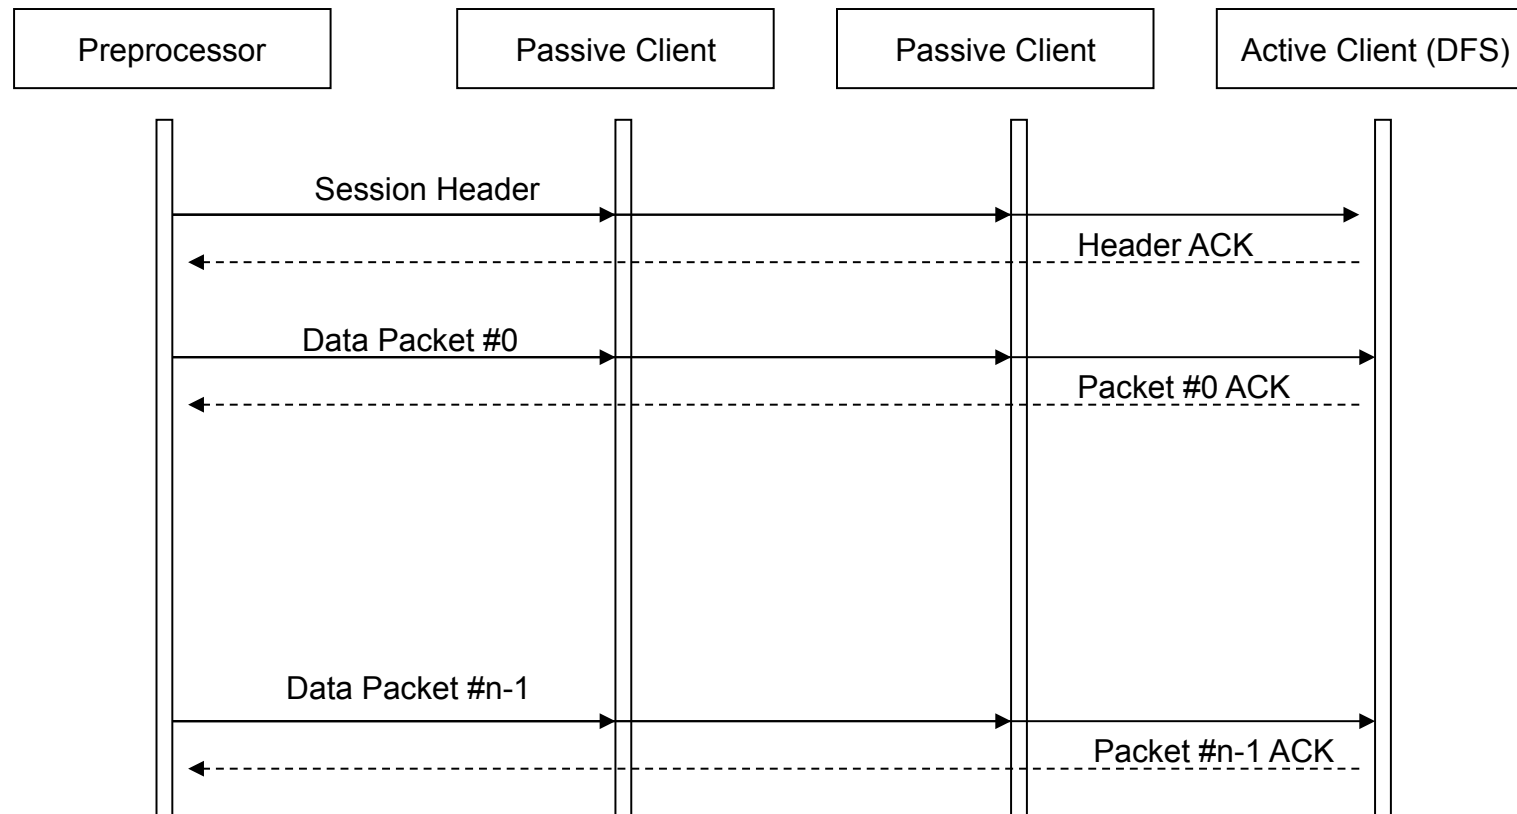


Data Link-Early Implementation



Broadcasting the Data Stream

Guaranteed Delivery To One Point



Shortcomings

- Same issues with CAsyncSocket.
- At data rates of 200K events/sec, end point dropped packets requiring clean up sends from preprocessor local backup.
- Required Waiting for ACK (poor bandwidth utilization).
- Session start protocol could be cleaner.
- PulselD info missing.



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Current Implementation

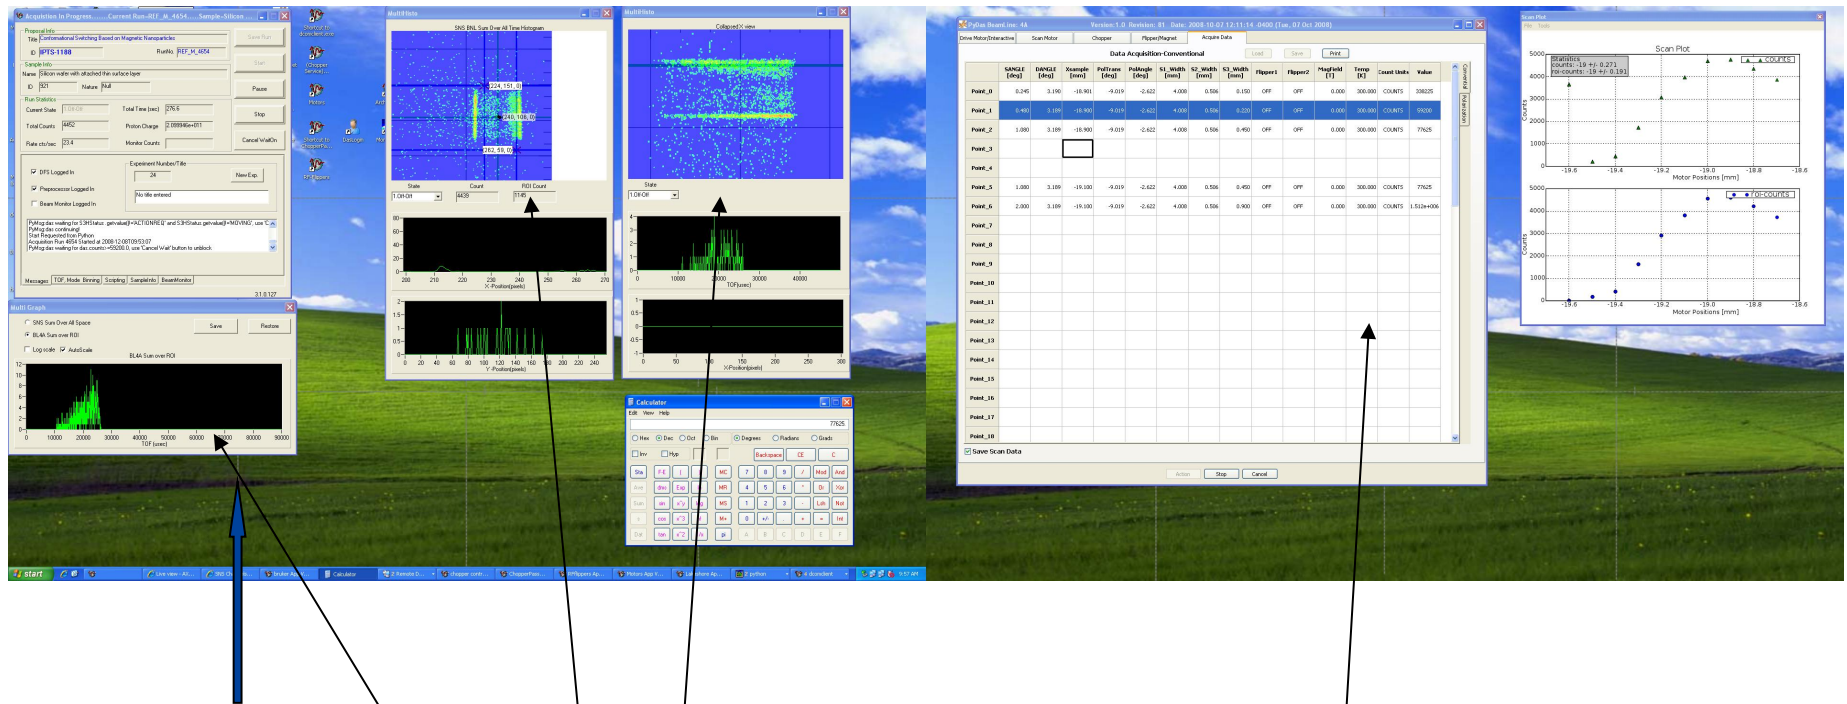
- TCP link using dedicated network to DataFileServer (DFS).
- UDP no longer waits for ACKS (true broadcast scheme).
- No “special” data packets in UDP stream and data broadcast contains PulseID info.
- DFS and Control Computer Listeners use multithreaded listening scheme



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



SNS Control and Monitoring GUIs



Main process has 8 to 10 threads running.

Each display window uses separate thread.

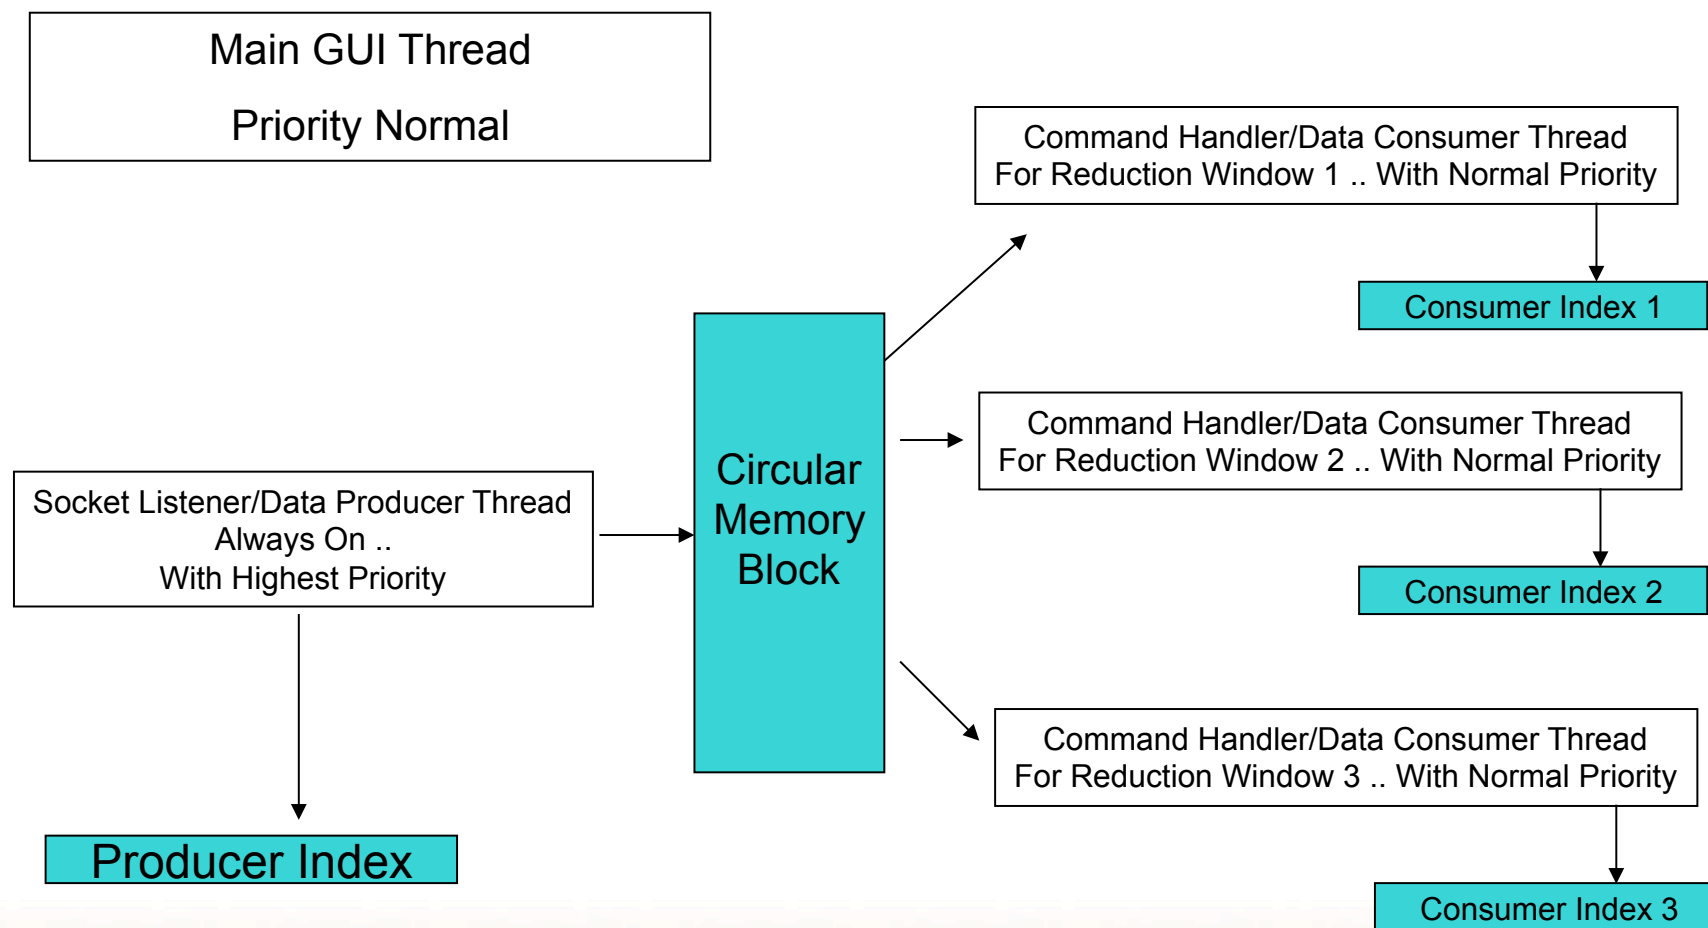


OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



UT-BATTELLE

Control Computer Data Link Threading Model. One Producer-n Consumers.



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Using Threads and Circular Buffers Elsewhere.

- One Producer-Three Consumers on Preprocessor.
- Very useful for logging. (I.e. strings are added to memory, separate thread writes to disk.)



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Benchmarks-Reliability

- TCP link 100% reliable.
- UDP link: 10,000 4K packets .75-1.1 sec 75% to 99.6% reliable.
10,000 8K packets .75-1.72 sec 86%- 99.6% reliable
10,000 16K packets 2.1-2.5 sec 98.35 -99.6% reliable

Still dropping packets.!!!


- UDP link with setsockopt(RCVBUF=1Meg). 100% reliable.



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Benchmarks-Performance



- With Jumbo Frames can obtain wire-speed transfers or 10^{-3} worse speeds. Packets must be larger than MTU.
- Without Jumbo Frames, UDP results similar to slightly better than TCP.

Note: NETPERF always gives poor results: 33% of wirespeed.



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY



Lessons Learned/Best Practice

- Take care using convenience classes.
- Thread whenever possible.
- Be aware of slow or blocking function calls in “performance” threads.
- TCP not always slower than UDP, best throughput always with Jumbo Frames.
- Read the fine print when using socket function calls.
- Don’t be surprised to see a few mountain goats.



OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

